



US 20090257350A1

(19) **United States**
(12) **Patent Application Publication**
Bugenhagen

(10) **Pub. No.: US 2009/0257350 A1**
(43) **Pub. Date: Oct. 15, 2009**

(54) **SYSTEM AND METHOD FOR USING NETWORK PERFORMANCE INFORMATION TO DETERMINE IMPROVED MEASURES OF PATH STATES**

Related U.S. Application Data

(60) Provisional application No. 61/047,028, filed on Apr. 22, 2008, provisional application No. 61/043,551, filed on Apr. 9, 2008.

(75) Inventor: **Michael Bugenhagen**, Olathe, KS (US)

Publication Classification

(51) **Int. Cl.**
H04L 12/24 (2006.01)
H04L 12/26 (2006.01)

Correspondence Address:
SONNENSCHN NATH & ROSENTHAL LLP
P.O. BOX 061080, WACKER DRIVE STATION,
WILLIS TOWER
CHICAGO, IL 60606-1080 (US)

(52) **U.S. Cl.** **370/235; 370/248**

(57) **ABSTRACT**

A method of determining the state of a path in a network includes measuring at a node of the network a rate of throughput over an interval for at least a portion of the path and determining at the node of the network a derivation associated with a characteristic of the frames communicated over at least the portion of the path. The method also includes determining a state of congestion of the path in response to the measured rate of throughput and the determined associated derivation.

(73) Assignee: **EMBARQ HOLDINGS COMPANY, LLC**

(21) Appl. No.: **12/421,525**

(22) Filed: **Apr. 9, 2009**




FIG. 1




FIG. 2




FIG. 3

300

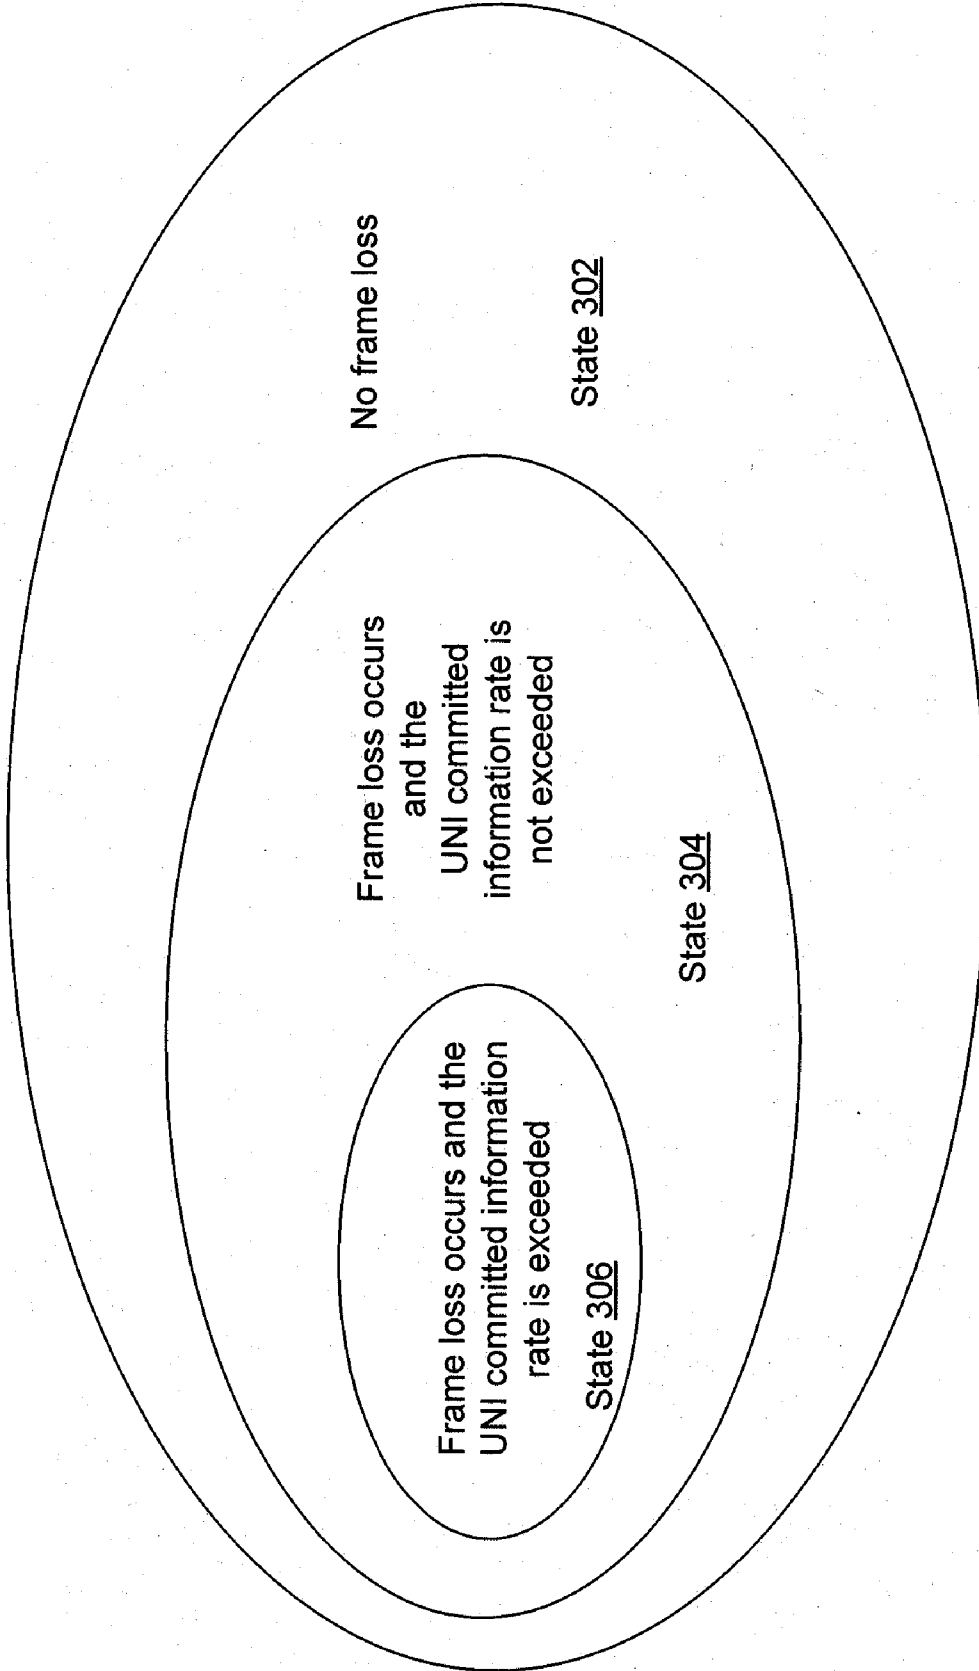


FIG. 4

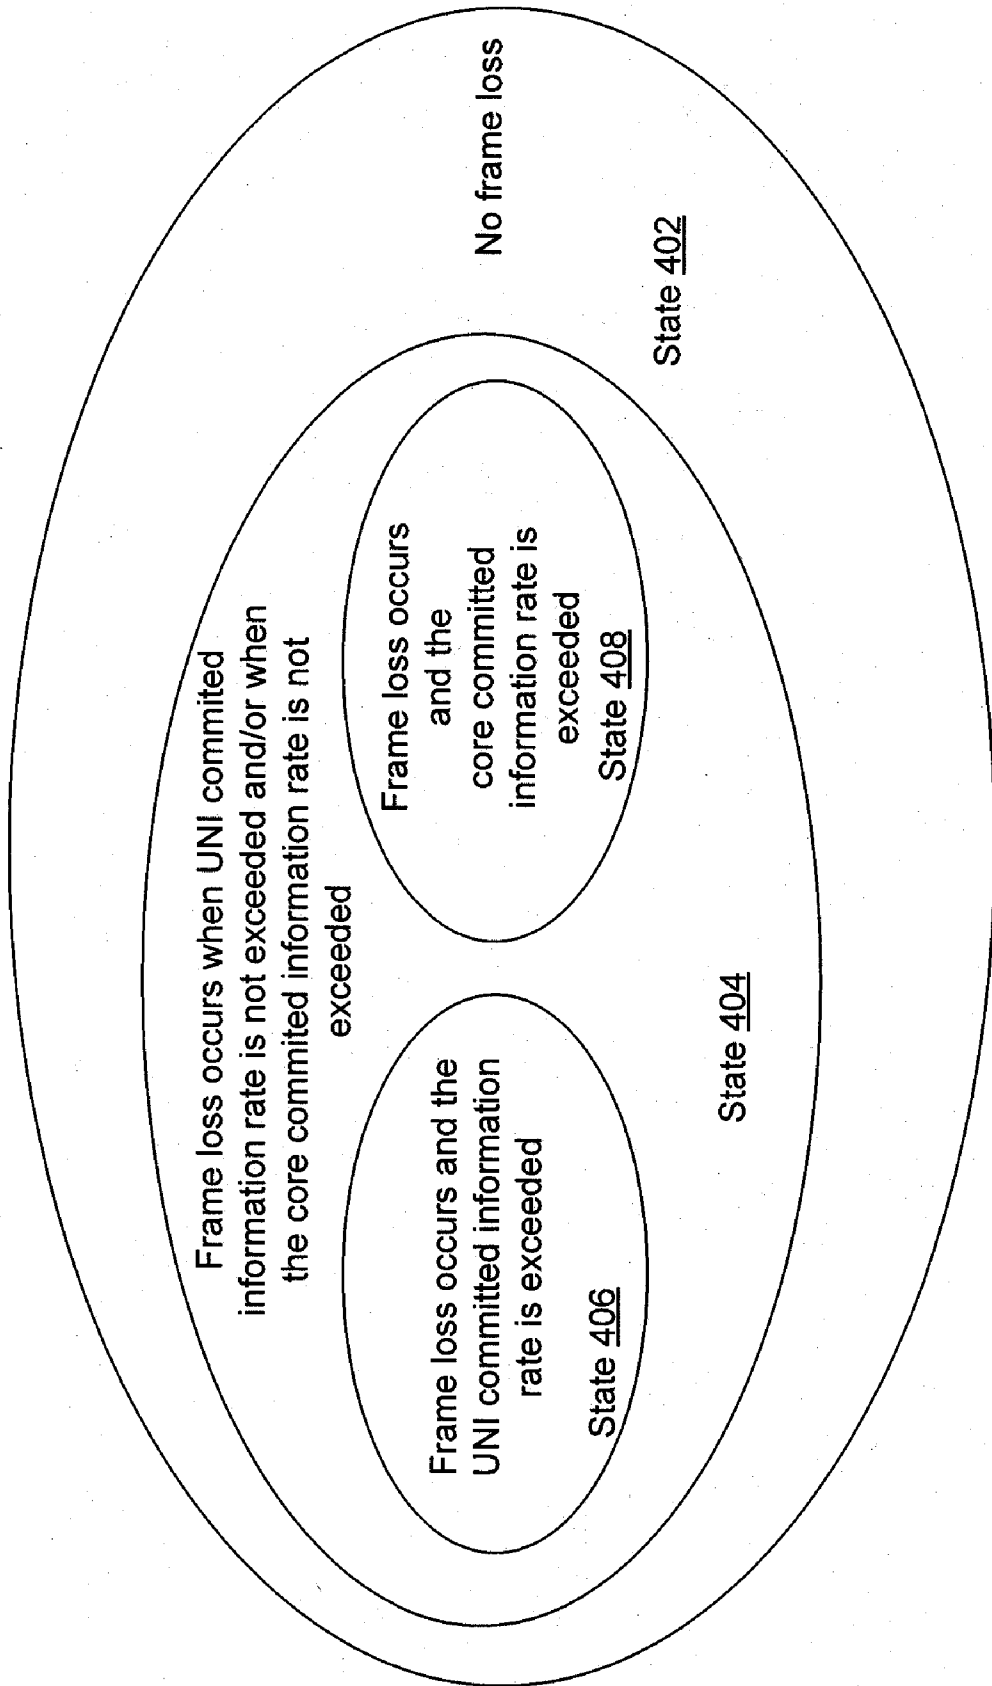


FIG. 5

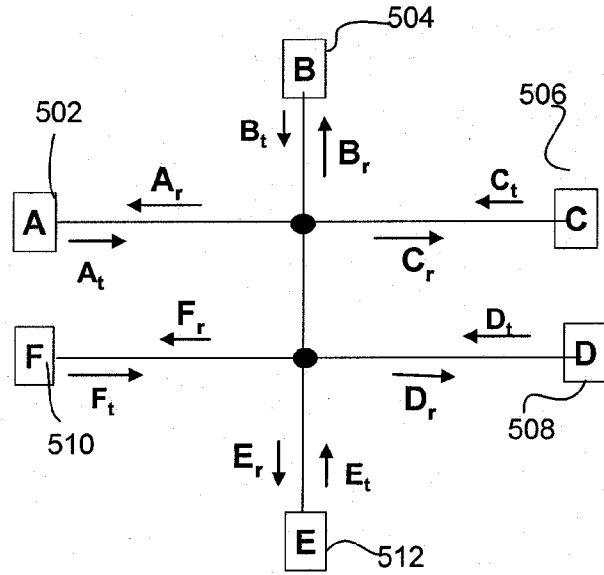


FIG. 6

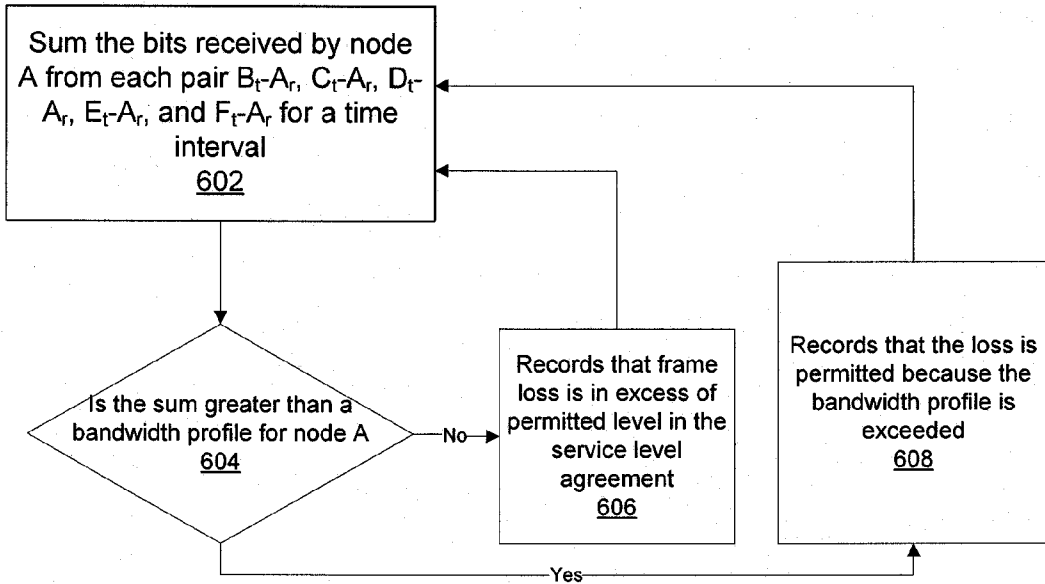


FIG. 7

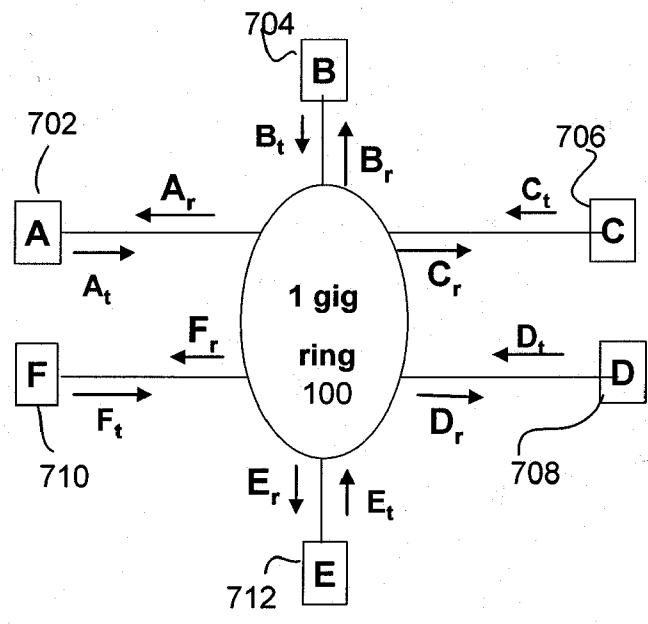


FIG. 8

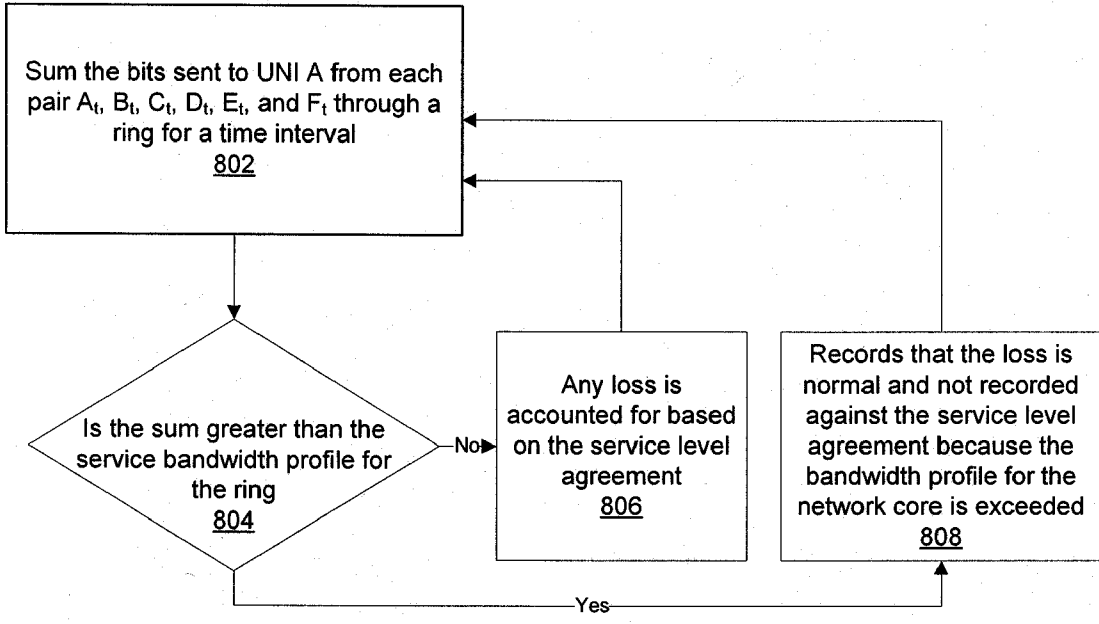



FIG. 9



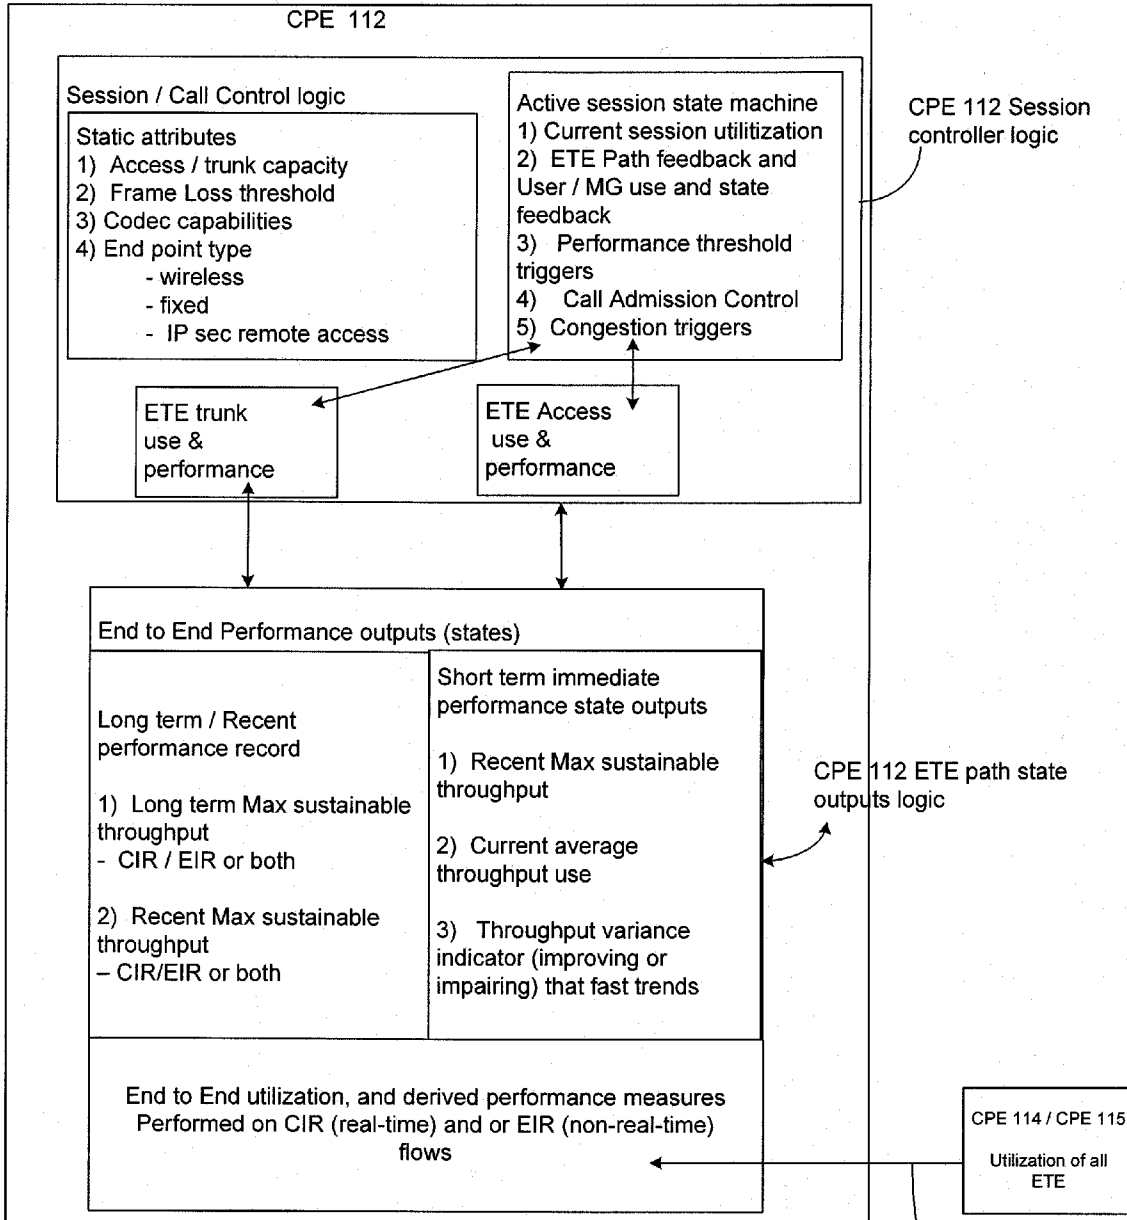


FIG. 10

Packet path ETE performance via packet path (optionally via external EMS)

SYSTEM AND METHOD FOR USING NETWORK PERFORMANCE INFORMATION TO DETERMINE IMPROVED MEASURES OF PATH STATES

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to provisional applications: Ser. Nos. 61/047,028 and 61/043,551 filed on Apr. 22, 2008 and Apr. 9, 2008, and U.S. patent application Ser. No. 12/381,673 filed on Mar. 12, 2009, entitled: SYSTEM AND METHOD FOR TRACKING PERFORMANCE AND SERVICE LEVEL AGREEMENT COMPLIANCE FOR MULTIPOINT PACKET SERVICES, respectively, which are incorporated herein by reference.

BACKGROUND

[0002] The use of and development of communications has grown nearly exponentially in recent years. The growth is fueled by larger networks with more reliable protocols and better communications hardware available to service providers and consumers. In particular, an Ethernet local area network (E-LAN) service type may be used to create a broad range of services. E-LAN service types may be utilized based on one or more service level agreements (SLA) for a multipoint service.

[0003] In many cases, the applicable communications network may include any number of service providers, access providers, legs, customers, and other elements that may complicate tracking performance or compliance for users or customers. The performance metrics are useful for trouble shooting, fault isolation, performance management (PM) threshold crossing, error identification, and other measurements that may not be shared between the many separate parties.

SUMMARY

[0004] One embodiment of the present invention is a method of determining the state of a path in a network. The method includes measuring at a node of the network a rate of throughput over an interval for at least a portion of the path and determining at the node of the network a derivation associated with a characteristic of the frames communicated over at least the portion of the path. The method also includes determining a state of congestion of the path in response to the measured rate of throughput and the determined associated derivation.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Illustrative embodiments of the present invention are described in detail below with reference to the attached drawing figures, which are incorporated by reference herein and wherein:

[0006] FIG. 1 is a pictorial representation of a communications system implemented in accordance with an illustrative embodiment;

[0007] FIG. 2 is a pictorial representation of an E-LAN service type utilizing multipoint-to-multipoint Ethernet virtual connections (EVCs) in accordance with an illustrative embodiment;

[0008] FIG. 3 is a state diagram of SLA states in accordance with an illustrative embodiment;

[0009] FIG. 4 is a state diagram of SLA states in accordance with an illustrative embodiment;

[0010] FIG. 5 is a pictorial representation of nodes in a network in accordance with an illustrative embodiment;

[0011] FIG. 6 is a pictorial representation of a flowchart of a process for determining whether losses within a network comply with a SLA in accordance with an illustrative embodiment;

[0012] FIG. 7 is a pictorial representation of a network ring in accordance with an illustrative embodiment;

[0013] FIG. 8 is a pictorial representation of a flowchart of a process for determining whether losses within a network comply with a SLA for core throughput in accordance with an illustrative embodiment;

[0014] FIG. 9 is a pictorial representation of a network performance system in accordance with an illustrative embodiment.

[0015] FIG. 10 is a pictorial representation of a CPE system in accordance with an illustrative embodiment.

DETAILED DESCRIPTION OF THE DRAWINGS

[0016] Certain illustrative embodiments of the described inventive concepts may be embodied in hardware, software, firmware, or a combination thereof. For example, in one embodiment the processes described herein may be encoded on a computer-readable medium, a hard disk drive, an ASIC, a PGA, ROM or any other suitable storage device or digital circuitry suitable to be utilized with or within a network device. In one embodiment, bandwidth measurements and other performance metrics may be monitored by logic engines, devices, or other elements to suppress alarms during specific conditions when the alarms are not intended to be activated.

[0017] Information regarding performance, state, and SLA compliance may be utilized by one or more embodiments of the present invention to determine that rate limiters or rate shapers have not sufficiently limited data traffic in a system to account for potential situations or inherent limitations. As a result, line rates have exceeded the sustainable throughput levels of the applicable ports or devices. Certain illustrative embodiments may utilize a static bandwidth profile or a dynamic bandwidth profile that accounts for different frame sizes to determine the thresholds and parameters that may be utilized to determine whether the SLA is compliant.

[0018] Two “normal” congestion states may occur in a multipoint service. The first congestion state is congestion experienced through the UNI (User Network Interface) egress point (hereafter “egress congestion” or “UNI congestion”) and the second is the congestion experienced with the core network service allocated bandwidth. UNI egress congestion may occur within multipoint services because the ingress of each port is “over-subscribed” based on the fact that multiple end-points (other UNI’s) may coincidentally transmit traffic to a single destination UNI at one given time. In other words, multiple UNIs may communicate with a single UNI simultaneously. This creates an increased likelihood of congestion for multipoint services. As a result of such congestion at such a UNI egress point, point-to-point performance data may incorrectly show congestion between two points despite the fact that there is no congestion on the remainder of network services. To correctly characterize the multipoint service for purposes of reporting service level performance to customers, the performance information regarding congestion at any UNI egress point during the congestion period should be removed for more accurate SLA tracking.

[0019] The second congestion state is congestion experienced through a segment of a network used by more than one customer, such as a segment of a core network (hereafter “core congestion”). This may occur with special services and is useful when the SLA contains a “backbone” or global bandwidth limitation. In this case, the ability to transmit from all UNI ports at one time is restricted to a global amount of bits per second. This limitation may introduce a known and expected congestion state for the service. To effectively track the multipoint service using point-to-point frame loss or other performance criteria, the “congestion state” of the backbone network may need to be independently tracked. Similarly, the performance metrics for the backbone network may need to be removed from the SLA metrics to effectively characterize the network service being offered according to a particular SLA.

[0020] In one embodiment, a method may be embodied in a set of instructions stored within a memory and executed to determine compliance or non-compliance with a particular SLA. Certain illustrative embodiments may be utilized for unicast or multicast traffic. Certain illustrative embodiments may be utilized to track the state of performance and devices within communications networks and systems. In one embodiment, certain illustrative embodiments may be used to resolve SLA issues between service providers, operators, customers, and other users. In particular, certain illustrative embodiments may be utilized to display, track, and store information, data, and empirical evidence of SLA compliance and noncompliance across any number of parties, devices, or elements. In one embodiment, the use of counters and performance measurement functions may track all frames regardless of priority markings. In a second embodiment, the counters and performance measurement functions may track the performance of the Committed Information Rate (CIR) frames. In a third embodiment, counters and performance measurement functions may track the Extended Information Rate (EIR) performance. In the fourth embodiment, the counters and performance measurement functions may track both CIR and EIR frames. Each embodiment may utilize a specific set of counters that measure only frames with specific EVC, P-bit, and/or other layer 3 (or other layer) marking methods to identify the traffic as EIR or CIR. As used herein, the term frame may be used to refer to frames, cells, packets, or any other suitable subset of data communicated over a network.

[0021] FIG. 1 is a pictorial representation of a communications system implemented in accordance with an illustrative embodiment. The communication system **100** of FIG. 1 includes various elements that may be used for wireless and wired communication. The communications system **100** may include a communications network **102**, a communications management system **104**, a server **106**, UNIs **108**, **110**, and **111**, customer premise equipment (CPEs) **112**, **114**, and **115**, and intermediary devices **116** and **118**. The communications system **100** may include any number of these elements, devices, components, systems, and equipment in addition to other computing and communications devices not specifically described herein for purposes of simplicity. For example, the communications system **100** may include various rate limiters or bit shapers. The different elements and components of the communications system **100** may communicate using wireless communications, such as satellite connections, WiFi, WiMAX, CDMA wireless networks, and/or

hardwired connections, such as fiber optics, T1, cable, DSL, high speed trunks, and telephone lines.

[0022] Communications within the communications system **100** may occur on any number of networks which may include wireless networks, data or packet networks, cable networks, satellite networks, private networks, publicly switched telephone networks (PSTN), communications network **102**, or other types of communication networks. The communications network **102** is an infrastructure for sending and receiving messages and signals according to one or more designated formats, standards, and protocols. The networks of the communications system **100** may represent a single communication service provider or multiple communications services providers. The features, services, and processes of the illustrative embodiments may be implemented by one or more elements of the communications system **100** independently or as a networked implementation.

[0023] In one embodiment, the communications network **102** is a Metro Ethernet network (MEN). Metro Ethernet network is a computer network based on the Ethernet standard and covering a metropolitan area. A Metro Ethernet network may be used as a metropolitan access network to connect subscribers and businesses to a wider area network, such as the Internet. Certain illustrative embodiments may be implemented utilizing any number of packet based networks or services, such as E-LAN or VLAN.

[0024] In one embodiment, an E-LAN service type may provide a best effort service with no performance assurances between the UNIs **108**, **110** and **111**. The UNIs are the physical and electrical demarcation point between the user and the public network service provider. In one embodiment, the UNIs **108**, **110** and **111** connect a MEN from the CPEs **112**, **114** and **115**.

[0025] In another embodiment, the communications service or E-LAN service type may include performance assurances, characteristics, thresholds, parameters, and information between UNIs **108**, **110** and **111**, such as a committed information rate (CIR). CIR is a specified amount of guaranteed bandwidth. Other performance information may include a committed burst size (CBS), an excess information rate (EIR) with an associated excess burst size (EBS), delay, delay variation, loss, and availability for a given class of service (CoS) instance. For example, EIR may be a throughput performance management that tracks all transmitted and received frame octets.

[0026] In another example, CIR may be a service performance management of conforming traffic that represents the frame loss threshold used to determine if the service is conforming to the SLA.

[0027] In one embodiment, the performance guarantees are included in an SAA. The SLA specifies the communications guarantees, thresholds, and actions that are agreed to by the communications service provider and a user/customer. Each of the UNIs **108**, **110** and **111** may have a specified bandwidth CIR.

[0028] Configurations such as multipoint communications may introduce natural conditions, such as oversubscription. Bandwidth congestion states may result under and SLA when multiple UNIs communicate with a single UNI simultaneously. Frame loss may be fully acceptable when a UNI is at the specified CIR, indicating the user or customer is breaking the boundary or threshold of allowable service.

[0029] In one embodiment, the CPEs **112**, **114** and **115** may be routers. In another embodiment, the UNIs **108**, **110** and

111 may be switches or other intelligent network devices. The UNIs **108**, **110** and **111**, the CPEs **112**, **114**, and **115**, the server **106** and other computing and communications devices within the communications network system **100**, which may include busses, motherboards, circuits, ports, interfaces, cards, connections, leads, transceivers, displays, antennas, and other similar components. The UNIs **108**, **110** and **111**, the CPEs **112**, **114** and **115**, and the server **106** may further include a processor and memory as well as other communications and computing elements including, but not limited to busses, circuits, cards, boards, ports, caches, buffers, power supplies, drives, and other components. In one embodiment, certain illustrative embodiments may be implemented by instructions stored within the memory. In another embodiment, the logic may be integrated, programmed, or stored within a circuit, chip, or card.

[0030] The processor is circuitry or logic enabled to control execution of a set of instructions. The processor may be a microprocessor, digital signal processor, central processing unit, application specific integrated circuit, or other device suitable for controlling an electronic device including one or more hardware and software elements, executing software, instructions, programs and applications, converting and processing signals and information, and performing other related tasks. The processor may be a single chip or integrated with other computing or communications elements.

[0031] The memory is a hardware element, device, or recording media configured to store data for subsequent retrieval or access at a later time. The memory may be static or dynamic memory. The memory may include a hard disk, random access memory, cache, removable media drive, mass storage, or configuration suitable as storage for data, instructions, and information. In one embodiment, the memory and processor may be integrated. The memory may use any type of volatile or non-volatile storage techniques and mediums. In one embodiment, the memory may store the performance management information, data, and states that are determined and tracked as herein described. The memory may include any number of databases for tracking transmitted and received packets from one or more UNIs, nodes, maintenance entities, or other devices, elements, or modules.

[0032] As described, the UNIs **108**, **110** and **111** may determine whether the terms of an SLA are being met utilizing congestion states, dynamic bandwidths, and throughput comparisons. The UNI congestion and core congestion may be utilized to analyze SLA compliance or noncompliance as well as generation of relevant alarms.

[0033] FIG. 2 is a pictorial representation of a communication system utilizing multipoint-to-multipoint Ethernet virtual connections (EVCs) in accordance with an illustrative embodiment. The communication system **100** is one embodiment of an E-LAN service type. The communication system **200** may include a MEN **201**, and UNIs **202**, **204**, **206**, and **208**.

[0034] In various embodiments, such as the multipoint configuration of FIG. 1, measuring the CIR performance may be difficult. The difficulty may result from communications or inter-connectivity between the UNIs **202**, **204**, **206**, and **208**. UNIs **202**, **204**, **206**, and **208** function as or are matched with one or more customer side UNI (UNI-C) and network side UNIs (UNI-N). The communication system may further include one or more network to network interfaces (NNIs), which may include external NNIs (E-NNI) and internal NNIs (I-NNI). Performance measurements and determinations of

the illustrative embodiments may be performed by frame, bits, or packet counters or other devices at the port level of one or more communications devices within the communication system **100**, which may also include the UNIs **202**, **204**, **206**, and **208**.

[0035] The communication system **100** illustrates one example of a shared bandwidth service in which the UNIs **202**, **204**, **206**, and **208** may contend for bandwidth to communicate with each other. The communications service provider operating the MEN **201** may define a known bandwidth profile for each of the UNIs **202**, **204**, **206**, and **208** which may include a bandwidth commitment. The bandwidth commitment may not be statically assigned to any specific UNI pair. As a result, using a single or even multiple UNI-to-UNI CIR bandwidth performance measurements may not be adequate to measure the parameters of the SLA agreed to by the communications service provider and the user.

[0036] In some cases, the communications service provider may guarantee bandwidth between legs of the service. The legs are connections between one or more UNIs. A guarantee between any two legs of the services reflects a core CIR. The SLA may provide guarantees for the core network, individual legs, and other point to point communications. In order to conform with the SLA, any of the UNIs **202**, **204**, **206**, and **208** may be required to transmit (i.e. tx) and/or receive (i.e. rx) a throughput CIR. In some cases traffic contentions (i.e. multiple UNIs communicating with a single UNI simultaneously) may result in frame loss because of the contention for the shared bandwidth. If a frame loss ratio is utilized as a SLA indicator, the amount of bandwidth frames that may be communicated without contention is significantly smaller than the full bandwidth profile. Certain illustrative embodiments provide a system and method for addressing contention problems to individual UNIs, across legs, or through the network.

[0037] In one example, a communications service provider may sell service to four customers as represented by the UNIs **202**, **204**, **206**, and **208**. The SLA with each of the customers may guarantee an amount of bandwidth to each customer UNI, such as 1 Gb/second. However, problems result when the customers are communicating with one another at the same time. A SLA measurement for monitoring is performed in order to determine compliance. Communication between customers occurs frequently and could be web traffic, streaming data, phone calls, or other forms of communications.

[0038] In one example, the SLA guarantee for the four node configuration shown in FIG. 2 may be determined by taking the CIR throughput level to the customer represented by UNI **202**, 1 Gig, and dividing by three (the customer cannot be in communication with itself) to determine the guaranteed CIR is approximately 333 Mb.

[0039] Based on new interconnection standards and network configurations using frame loss as a metric to determine an SLA guaranteed CIR is problematic because of potential communications between the UNIs **202**, **204**, **206**, and **208**. Similarly, testing SLA guarantees via the application of throughput stress testing to confirm a CIR may cause problems by breaking operational rules, disrupting live service traffic to any number of other customers, and potentially crashing applications or portions of the MEN **201**. Typically, stress testing degrades the shared bandwidth of the service itself and may not capture transient defective states that may occur. An illustrative embodiment may be utilized to track throughput performance measurements that may alleviate the

need for in-service traffic testing utilizing synthetic or test packets that do not represent normal network traffic.

[0040] In one embodiment, in order to set a CIR frame loss rate level so that it may not indicate loss, the level may be set to a threshold in which no contention occurs on the UNI. For example, $CIR = UNI \text{ bandwidth} / (\# \text{ nodes} - 1)$. In another embodiment, the CIR may be set so that it is no more than fifty percent of the service offered to the UNIs **202**, **204**, **206**, and **208**.

[0041] The bandwidth utilization or CIR utilization may be measured numerically or statistically for each potential communications path in the MEN for communication data. For example, throughput counters and frame loss indicators may monitor and record relevant information and performance metrics. The traffic may be measured utilizing any number of devices or applications that track, measure, and/or record network traffic measurements. Frame loss is only expected when the SLA limits are exceeded or UNI congestion occurs. For example, the UNI **206** may have a CIR of 100 Mb. No frame loss is expected if less than 100 Mb are received, if however, UNI **202** transmits 80 MB to UNI **206**, UNI **204** broadcasts 40 Mb, and UNI **208** broadcasts 30 MB, the incoming bandwidth is 150 Mb exceeding the CIR of 100 Mb. As a result, frame loss is expected and the SLA is considered to be compliant despite the frame loss.

[0042] Second, the core of the MEN **201** itself may have a core CIR. Tracking frame loss rate between all UNIs **202**, **204**, **206**, and **208** does indicate if congestion or faults are occurring at a UNI if the core is has reached an SLA capacity. During the time frames when the service is running at full capacity frame loss rate becomes expected or is within SLA compliance or acceptable behavior. For example, the MEN **201** may have a core CIR of 800 Mb. If UNIs **202**, **204**, and **206** communicate 250 Mb and UNI **208** communicates at 200 Mb, the core CIR of 800 Mb is exceeded by 150 Mb and frame loss is determined to be acceptable. As a result, potential alarms are not generated and the communications service provider is determined to have not violated the SLA despite the frame loss. Tracking UNI congestion and core congestion through frame loss indicators and UNI states enable true SLA conformance measurements.

[0043] The examples given of 95% of CIR and 1% are illustrative thresholds, UNI 1 tracks TX, RX for itself (i.e., alarms on the RX of UNI 1). The measurements and calculation of throughput and frame loss may be accomplished using any suitable packet counter, network probe, analysis tool, which may be integrated with a UNI, router or other network node or be displayed in line with such a device.

[0044] FIG. 3 is a state diagram of SLA states in accordance with an illustrative embodiment. The state diagram **300** of FIG. 3 may be embodied or implemented in a chip or chipset, digital logic, fully programmable gate arrays, or an application for determining the status of an SLA. In one embodiment, FIG. 3 may be applicable to Scenario 1 described above. The state diagram **300** may include states **302**, **304**, and **306**. Service providers, operators, customers, and other groups may need the ability to isolate states of the network. The states may be utilized to minimize service interruptions, repair times, and operational resources by detecting, diagnosing, localizing, and notifying network management systems of defects in order to take corrective actions appropriate to the type of defect.

[0045] The state diagrams of FIGS. 3 and 4 may be utilized at one or more nodes in a network to determine whether the

SLA is being met and record live traffic throughput to capture applicable thresholds as frame loss begins to occur. In one embodiment, the receiving portion of the UNI or node may perform all determinations of SLA compliance. In another embodiment, the applicable states may be performed by a measurement information base (MIB). States **302**, **304**, and **306** may indicate an alarm state, suppress alarms or take one or more associated actions. The states may be utilized to dynamically determine usage and oversubscription information and profiles for traffic engineering and network management and design. The states may also be utilized to record and store network or UNI characteristics and performance information at the moment(s) the alarm state is triggered.

[0046] In state **302** and state **306**, the SLA is in compliance. In state **304**, the SLA is in non-compliance. In state **302**, no frame loss occurs. The SLA may be considered to be compliant.

[0047] In state **304**, frame loss occurs in excess of the frame loss permitted under the SLA and the UNI CIR is not exceeded. As a result, the SLA is considered to be non-compliant. The SLA is non-compliant because the CIR is not being exceeded by the customer, but yet there is still frame loss in excess of the frame loss permitted under the SLA that does not fall within the terms of the SLA. As a result, any number of alarms may be activated or asserted indicating frame loss for troubleshooting, diagnose, and other corrective actions.

[0048] In state **306**, frame loss occurs in excess of the frame loss permitted under the SLA and at the same time, the UNI CIR is exceeded. State **306** is acceptable because the user has exceeded the bandwidth guaranteed to the user at the transmit or receive end of the customer port and as a result frame loss and non-compliance with the SLA is acceptable. During state **306**, any number of alarms that may be activated due to frame loss may be ended or turned off because the UNI CIR is exceeded. The determinations of the alarm state utilizing states **302**, **304**, and **306** may be performed locally by a device or performance information may be communicated to a separate network device for determination and alarm state management and control.

[0049] FIG. 4 is a state diagram of SLA states in accordance with an illustrative embodiment. The state diagram of FIG. 4 may be implemented or function with other state diagrams, such as FIG. 3. The state diagram **400** may include states **402**, **404**, **406**, and **408**. In one embodiment, FIG. 4 may be applicable to Scenario 2 described above. In some cases, service providers do not provide broad guarantees for bandwidth. For example, the communications service provider may limit the SLA based on a leg CIR and a core CIR through the network. As previously described, in the state **402** no frame loss occurs in excess of the frame loss permitted under the SLA and the SLA is considered in compliance.

[0050] In state **404**, frame loss occurs when the UNI CIR is not exceeded and/or when the core CIR is not exceeded. As a result, the SLA is considered noncompliant in state **404**. Any number of alarms or indicators may be set or initiated in response to entering state **404**.

[0051] In state **408**, frame loss occurs in excess of the frame loss permitted under the SLA when the core CIR is exceeded, the SLA is determined to still be in compliance. In state **406**, if frame loss occurs when the UNI CIR is exceeded, the SLA is determined to still be in compliance. During states **406** and

408, alarms, such as those activated for state **404**, may be deactivated, cancelled, or disengaged because of the noncompliance with the SLA.

[0052] In one embodiment, the communications service provider may track the packets transmitted and received over the core network (all UNIs bundled together) between each UNI pair, and to each and from each UNI.

[0053] For SLA compliance, a throughput SLA does not provide accurate information on SLA compliance. As a result, frame loss is still required as an important metric to determine compliance with the SLA requirements. Determining SLA compliance may be performed as described by the various illustrative embodiments.

[0054] Certain illustrative embodiments may allow a communications service provider or customer to measure metrics used to determine whether obligations are being met under an SLA. The proposed systems and methods do not require stress testing the network in order to determine SLA compliance. Additionally, certain illustrative embodiments may be utilized by network engineers to determine potential traffic and demonstrate compliance with SLAs when customers are operating within their CIRs. Additionally, communications service providers may use Certain illustrative embodiments to ensure that they do not oversell bandwidth, legs, or the core user to provide their service.

Bandwidth

[0055] Throughput performance management rates are frequently discussed in terms of Ethernet frame octets received (EFOR) and Ethernet frame octets transmitted (EFOT). The throughput rate is the amount of EFOR and EFOT measured over a short time period. In one embodiment, the throughput rate may be measured with respect to the bits passing through a maintenance entity, network node, or Access Transport Resource Control Functional (A-TRCF) entity for the CIR, extended information rate (EIR), or both together as a single measure of throughput. Performance measurements, such as frame loss, may be calculated based on all network traffic or only based on conforming traffic, such as traffic that is within a customer CIR.

[0056] Ethernet has multiple standard frame sizes whereas asynchronous transfer mode (ATM) has one frame size. In one example, Ethernet frames may vary in size from 64 to 1,518 octets. This however, does not include 96,200 jumbo frames supported by Gig-E. Live Ethernet traffic includes mixed types of packet sizes with Voice over Internet Protocol (VoIP) packets generally being around 128 bytes, and Internet traffic being composed of both 64 byte and 1,518 byte packets. Given that the service blend of packets of differing size is dynamic and that differences may exceed two orders of magnitude, such frame loss as a performance indicator is inaccurate. The deviation in packet size makes using frames per second or frame rate measurements an invalid throughput indicator or performance measurement. Clock skew in Ethernet chips may also cause variations as high as 1% in the amount of frames that may be transmitted due to variations in frame gaps. Although not illustrated herein, the embodiments described in FIGS. 3 and 4 may be expanded to further consider whether or not the CIR of a remote node or UNI transmitting to the illustrated UNI has been exceeded. If the CIR for such remote UNI have been exceeded, network performance measurements indicating a performance issue such as excessive frame loss that include the measurement of traffic to or from the remote UNI may also be permitted despite vio-

lating a particular service level because of the exceeded CIR. Although not illustrated herein, the embodiments described in FIGS. 3 and 4 may be expanded to further consider whether or not the CIR of a remote node or UNI transmitting to the illustrated UNI has been exceeded. If the CIR for such remote UNI have been exceeded, network performance measurements indicating a performance issue such as excessive frame loss that include the measurement of traffic to or from the remote UNI may also be permitted despite violating a particular service level because of the exceeded CIR.

[0057] FIG. 5 is a pictorial representation of nodes in a network in accordance with an illustrative embodiment. FIG. 5 shows multiple nodes that may be included in a network. In one embodiment, the network may include node A **502**, node B **504**, node C **506**, node D **508**, node E **512**, and node F **510**. As previously described, the nodes may represent any number of UNIs, devices, components, equipment, ports, or constructs.

[0058] The nodes may both transmit and receive data to any of the other nodes as shown. As previously described, a network performance device or process may not be able to determine compliance with a SLA when multiple nodes communicate with a single node simultaneously. For example, node B **504**, node C **506**, and node D **508** transmit data to node E **512** at the same time utilizing a large amount of bandwidth that exceeds the terms of the SLA. Certain illustrative embodiments provide a system and method for compensating for the situations of FIGS. 5 and 7 without generating alarms, recording SLA violations or entering an error state.

[0059] FIG. 6 is a pictorial representation of a flowchart of a process for determining whether losses within a network comply with an SLA in accordance with an illustrative embodiment. The process of FIG. 6 may be implemented by a UNI or other node in accordance with an illustrative embodiment. Although specifically described for one node in an illustrative embodiment, the process of FIG. 6 may be performed for multiple UNIs.

[0060] The process may begin with a UNI or other device summing the bits received by node A from each of nodes B through F for a specified time interval (step **602**). The nodes may represent UNIs communicating with the UNI being monitored, such as node A thereby establishing UNI pair Bt—Ar, Ct—Ar, Dt—Ar, Et—Ar, and Ft—Ar. The time interval may be any measure specified by the network administrator. In one embodiment, the time interval may be an amount less than 1 second. The measurement in step **602** measures the total amount of bits received by a node from all nodes within the network for such time interval. In one embodiment, the measurement or calculation of step **602** may be initiated in response to a determination that there is loss within the network or at the UNI being monitored.

[0061] Next, the UNI determines whether the sum of the summed bits is greater than a bandwidth profile for node A (step **604**). The bandwidth profile may be specified by the SLA. For example, the bandwidth profile for node A may be included in a customer agreement with a communications service provider.

[0062] If the UNI determines that the sum of the summed bits is greater than a bandwidth profile or CIR for node A, the UNI may indicate or store an indication that the frame loss triggers a violation of the SLA (step **606**). In one embodiment, the frame loss may trigger an obligation of the communications service provider. For example, the communications

service provider may be required to provide a customer associated with the UNI a discount, credits, or remedy for the loss.

[0063] If the UNI determines that the sum is greater than a bandwidth profile or CIT for node A, the UNI may indicate that the frame loss is permitted and does not trigger a violation of the service level agreement because the bandwidth profile has been exceeded (step **608**). Such indication may occur because multiple nodes are communicating with the UNI simultaneously in a manner that is not within a CIR or that is otherwise not supported by the SLA.

[0064] FIG. 7 is a pictorial representation of a network ring in accordance with an illustrative embodiment. FIG. 7 shows multiple nodes that may communicate through a ring **100** in a network. In one embodiment, the ring **100** may communicate with node A **702**, node B **704**, node C **706**, node D **708**, node E **712**, and node F **710**. As previously described, the nodes may represent any number of UNIs, devices, components, equipment, ports, or constructs. The ring **100** is the core of a network. The ring **100** is one embodiment of any number of topologies that may be utilized in a communications environment as all or a portion of a network. The ring **100** may include any number of interconnected devices, equipment, and systems. The nodes may communicate with the other nodes through the ring **100**, as shown. As previously described, the ring **100** may not be able to determine compliance with a SLA for guaranteed core performance when multiple nodes communicate through the ring **100** at or near maximum capacity simultaneously. For example, if node B **704**, node C **706**, node D **708**, and node E **712** transmit and receive data at their maximum bandwidth simultaneously, the terms of the core SLA may be exceeded. Certain illustrative embodiments provide a system and method for monitoring such situations without entering an error state.

[0065] FIG. 8 is a pictorial representation of a flowchart of a process for determining whether losses within a network comply with an SLA for core throughput in accordance with an illustrative embodiment. The process of FIG. 8 may be implemented by a UNI, application engine within a server, or other element of a communications management system. Although specifically described for one node, in an illustrative embodiment, the process of FIG. 8 may be performed for multiple nodes or UNIs simultaneously.

[0066] The process may begin by summing the bits received by UNI A from each node B_r , C_r , D_r , E_r , and F_r comprising five nod pairs through a ring over a particular time interval (step **802**). The pairs may represent all UNIs communicating directly with the UNI being monitored, such as UNI A. The time interval may be any measure specified by the network administrator. In one embodiment, the time interval may be an amount less than 5 seconds. The measurement in step **802** measures the total amount of bits received by UNI A from nodes within the network for the time period. In one embodiment, the measurement or calculation of step **802** may be initiated in response to a determination that there is frame or packet loss within the network or at the UNI being monitored, network congestion, or other network issues.

[0067] Next, the UNI determines whether the sum is greater than a bandwidth profile for the ring (step **804**). The bandwidth or throughput profile may be specified by the SLA. For example, UNI A may represent a UNI of a customer that has an agreement with a communications service provider for core bandwidth through a ring or core portion of the network.

[0068] If the UNI determines that the sum is not greater than a bandwidth profile for the ring, the UNI records that any

loss is accounted for based on the SLA. (step **806**). In one embodiment, the loss may be required to be accounted for by the communications service provider. For example, the communications service provider may be required to provide a customer associated with the UNI a discount, credits, or retribution for the loss.

[0069] If the UNI determines the sum is greater than a bandwidth profile for the ring, the UNI records that the loss is normal and not recorded against the SLA because the bandwidth profile for the network core is exceeded (step **808**). The recordation of step **808** indicates that the loss is occurring because multiple UNIs are communicating through the ring simultaneously in a manner that is unsustainable and not supported by the SLA. For example, the core bandwidth of the ring may be one gigabyte and may be exceeded by four UNIs simultaneously transmitting (or attempting to transmit) at 500 Mb.

[0070] Another matter complicating the issue of determining if a frame loss occurrence is normally expected is the use of a bit based rate limiter or shaper at one end of a path and a physical UNI port at the opposite end. Rate limiters enforce bandwidth profiles based upon bit rates and not frame sizes. UNI ports, however, have a frame size dependant bandwidth profile that limits the amount of "effective" bit throughput that may be transmitted through the port at any given time. This relationship is caused due to the "cell tax" or inter-frame-gap (IFG) and inter-frame-overhead, and other frame components, such as the start frame delimiter (SFD), that are not counted by bit based rate shapers. The end result of the UNI bandwidth profile is that for every frame per second transmitted through a UNI the "cell tax" overhead of the IFG, and other non-bits are subtracted from the UNI line rate. This relationship results in small frames having less "effective" throughput or bit based throughput than large frames. However, rate shapers and rate limiters do not change their bandwidth profile with frame size as do UNI ports. This causes a mismatch in the bit based rate limiter to frames size dependant throughput profile of a port. When a rate limiter uses a 76.2% or higher bandwidth profile of the UNI port line rate frame loss may occur because of the frame size dependency at the UNI port. To account for frame size, a frame per second dynamic bandwidth profile tool may be required to decipher if the bandwidth loss was caused by the bandwidth profile mismatch of a rate limiter to the UNI, which is frame size dependant.

[0071] FIG. 9 is a pictorial representation of a network performance system in accordance with an illustrative embodiment. In one embodiment, the network performance system **900** of FIG. 9 may be encompassed in a device, such as a UNI. A network performance system **900** may include or perform basic measurements, derived calculations, and threshold crossings. The elements may be circuits, logic, or hardware or may include a program, instructions, or stored elements. FIG. 9 may further include one or more interfaces communicating with a number of rate limiters or rate shapers. FIG. 9 may also be utilized to perform the other methods and processes herein described. In one embodiment, the network performance system **900** may include alarms **902**, performance logic **904**, SLA module **906**, dynamic bandwidth profile **908**, and logic **910** and **912**.

[0072] FIG. 10 shows one embodiment of end-to-end (ETE) path state information being received and or derived at a session or call controller represented by a CPE **112**. The ETE path state information may include network perfor-

mance information regarding an overall call or session connection path or a trunk path, access path, path between two UNIs, or any other desired path. Network performance information such as frame loss and actual or available throughput regarding an ETE path may determine an overall ETE path state based on network performance information regarding one or more communication links or network nodes included in the ETE path. More particularly, performance measurement functions may be utilized to measure network performance information and derive an ETE patent state at one or more network nodes, a central network resource such as an external EMS system, or at the session or call controller itself. A Network performance information may be communicated in-band or out-of-band. For example, to convey information between service providers and networks, in-band communication of network performance information and path state may be utilized. Information regarding historical ETE network performance and ETE path state may be stored and relayed to the higher application layers resident on a network element such as CPE 112. In one embodiment, CPE 112 in FIG. 1 is a call or session controller with a Media Gateway or other shared trunk device such as a pin-hole firewall, a CPE 114 is a gateway device, and a CPE 115 is a user device. In such an embodiment, the end-to-end path between CPE 112 and CPE 114 is a shared resource trunking packet path. In such embodiment, the path between CPE 112 and CPE 115 is an end-to-end access path. CPE 112 reads the end-to-end state information available regarding the path performance to make session setup, treatment, and processing decisions.

[0073] The end-to-end state information of the path may include one or more of the following: real-time traffic statistics normally marked as Committed Information Rate (CIR) traffic, non-real-time traffic statistics normally marked as Extended Information Rate (EIR) traffic, and any rejected traffic statistics. In one embodiment, both CIR and EIR are tracked together, in a second only CIR is tracked, and in a third both CIR and EIR are tracked separately.

[0074] The alarms 902 represent service or system alerts that may be generated in response to the rules of SLA module 906 being violated based on the dynamic bandwidth profile 904 as determined by the performance logic 904. The performance logic 904 is the circuits, instructions, and other intelligent components operable to determine whether the service level agreement has been violated utilizing a static bandwidth profile or the dynamic bandwidth profile 908. The SLA module 906 may utilize rules based on any number of states, steps, or processes as previously described in FIGS. 3-8.

[0075] The dynamic bandwidth threshold 908 may act to regulate a policing function by providing feedback that may be utilized to reduce throughput limits based on frames per second, or it may simply be used to gauge if a rate limiter failed to enforce a small frame bandwidth profile which was discarded by a UNI port. The dynamic bandwidth profile 908 may be applicable to ports or devices. The dynamic bandwidth profile 908 may utilize logic 910 or logic 912. In one embodiment, the dynamic bandwidth profile 908 may establish the maximum effective transmission rate based upon the UNI port's maximum frames per second at the largest maximum transmission unit (MTU) size, and the corresponding effective throughput rate for that MTU size and frame rate. In a second embodiment, the dynamic bandwidth profile may be based upon the UNI line rate, and the frames per second measure, with the corresponding frame overhead or "non-bit" counted portions of each frame. For example, the dynamic

bandwidth profile 908 may be utilized directly with the network bit based throughput counters to determine if the traffic present conforms to the port's frame based throughput instead of a static bandwidth profile in response to the dynamic bandwidth profile 908 which exists due to protocol overhead for small frames between 76.2% and 100% of the port line rate speed. The relevant speeds and percentages may be determined based on industry standard inter-frame-gap and "non-bit" overhead, such as SFD.

[0076] An example may be utilized to explain logic 910. In an illustrative embodiment, a single circuit with a Gig-E may be located at one end, and a 100 Mb Fast Ethernet port may be located at the second end. The service may be a 100 Mb service with a "rate-limiter" at the Gig-E end. In effect, the rate limiter will pass 100 Mbs of effective throughput. However, with small packets such as 64 byte frames only 76.2 Mbs of payload may fit inside of the 100 Mb line rate port due to overhead. However, 99.8 Mbs of 1581 byte frames will fit into the port due to a lower ratio of overhead to payload. The bandwidth profile of the rate limiter is not frame size dependant, and the bandwidth profile of the Fast Ethernet port is frame size dependant.

[0077] The rate limiter is typically a bit based effective throughput bandwidth profile may not be constrained by a line rate, and as a result may allow too many small frames through the service. As a result, the small frames may be dropped at the far end when the frames attempt to enter the UNI or 100 Mb Fast Ethernet port.

[0078] A dynamic frames per second "threshold" may be required to identify if the line rate of the frames being transmitted exceeded the bandwidth profile of the UNI port. Given that the frame bit counters commonly used by packet systems do not typically use a line rate, logic, an equation, algorithm, or other instructions may convert the instantaneous bandwidth (during a short time period) to a line rate that may be used to determine if the effective throughput may be communicated through the port or is less than the port speed.

[0079] The logic 912 may be utilized to convert the effective bandwidth to a line rate. If the effective bandwidth rate exceeds the UNI port line rate, the frame loss is not counted because those frames exceeded the UNI port rate (line rate, e.g., 100 Mbs). Normally, the effective throughput measurement does not detect such conditions.

[0080] The logic 912 may utilize a "static UNI line rate threshold" to determine if the throughput allowed by the rate limiter was or was not service conforming (able to be communicated through the UNI port). The logic 912 may utilize the instantaneous throughput determined by the bit counters and add the overhead of the IFG to the effective throughput to adjust to a line rate standard. That value is then compared to the line rate. If the value is above the port's line rate, the traffic does not conform to the service, if it was below the line rate (100 Mb), then the traffic did conform and the loss should not have occurred.

[0081] The following provides additional details and embodiments for the logic 910 and/or 912. In one embodiment, the dynamic bandwidth profile 908 may determine throughput rate of a port. The throughput rate may determined to be less than or equal to a determination of parameters. In one embodiment, the parameters may include effective throughput rate, port speed, the frames per second, and the inter frame overhead.

[0082] Throughput is dependent on frame size and may change as the frame size changes dynamically. As a result,

packet service level capacities are often stated in terms of the throughput rate of the MTU size or frame size. For example, 90 Mbs with 1518 byte frames. However, the use of smaller frames yields less throughput capacity. In one ongoing example, 64 bytes have a throughput capacity of 76.2 Megs for a 100 Megabyte port. In effect, every time a frame is added for a given time period, the size of the overhead and non-payload bits relative to the overall frame size also increases. For example, envelopes, preambles, headers, inter-frame gaps, trailers, OAM packets, performance information packets, and any other bits not associated with actual payload data would all be considered overhead or non-payload bits. Depending on the layer at which such throughput is measured, whether or not data is considered payload or non-payload bits may change. For example, data representing a header in one layer may be considered payload in another layer.

[0083] Therefore, the maximum MTU throughput rate minus the added incremental frames per second yields the instantaneous dynamic effective throughput capacity. Such logic may be expressed by subtracting the current frame per second interval from the fastest rate that frames may be sent per second for the maximum MTU size and then subtracted from the maximum MTU size effective throughput. The result is a threshold that indicates the amount of payload or effective bits that may pass through the service based on the loss of throughput due to the added frames per second overhead. In one embodiment, derivations may be made to determine average packet size, average header size, the number of payload bits, the number of non-payload bits, or any similar derivation based on one or more measurements or calculations. All of such derivations may be indicators used to determine a dynamic bandwidth profile **908** or similar indication of actual throughput or throughput capacity.

[0084] For example, the dynamic bandwidth profile **908** may be utilized to determine that the actual effective frame per second dependent bandwidth profile **910** is 76.2% of service payload utilizing 64 byte packets and 99.8% of service payload with 1518 byte packets. As a result, the systems, devices, user, or equipment may account for the potential 24% of overhead by adjusting line rates and throughput levels. The dynamic bandwidth profile **908** may ensure that frame loss does not occur between a UNI-N and a UNI-C.

[0085] FIG. 10 shows one embodiment of end-to-end (ETE) path state information being received and/or derived at a session or call controller represented by a CPE **112**. The ETE path state information may include network performance information regarding an overall call or session connection path or a trunk path, access path, path between two UNIs, or any other desired path. Network performance information such as frame loss and actual or available throughput regarding an ETE path may determine an overall ETE path state based on network performance information regarding one or more communication links or network nodes included in the ETE path. More particularly, performance measurement functions may be utilized to measure network performance information and derive an ETE patent state at one or more network nodes, a central network resource such as an external EMS system, or at the session or call controller itself a Network performance information may be communicated in-band or out-of-band. For example, to convey information between service providers and networks, in-band communication of network performance information and path state may be utilized. Information regarding historical ETE network performance and ETE path state may be stored and relayed to the higher application layers resident on a network element such as CPE **112**. In one embodiment, CPE **112** in

FIG. 1 is a call or session controller with a Media Gateway or other shared trunk device such as a pin-hole firewall, a CPE **114** is a gateway device, and a CPE **115** is a user device. In such an embodiment, the end-to-end path between CPE **112** and CPE **114** is a shared resource trunking packet path. In such embodiment, the path between CPE **112** and CPE **115** is an end-to-end access path. CPE **112** reads the end-to-end state information available regarding the path performance to make session setup, treatment, and processing decisions.

[0086] The end-to-end state information of the path may include one or more of the following: real-time traffic statistics normally marked as Committed Information Rate (CIR) traffic, non-real-time traffic statistics normally marked as Extended Information Rate (EIR) traffic, and any rejected traffic statistics. In one embodiment, both CIR and EIR are tracked together, in a second only CIR is tracked, and in a third both CIR and EIR are tracked separately.

[0087] FIG. 10 illustrates an active session state machine included within CPE **112**. The active session state machine may determine and store in memory a current session utilization, feedback on network performance information on the ETE path, feedback on an overall derived state on the ETE path. The active session state machine may include performance threshold triggers indicating triggers or levels whereby alarms, notifications, and changes are initiated relative to detected performance thresholds. The active session state machine may include guidelines for call admission control and call throttling. The active session state machine may include congestion triggers indicating triggers or levels whereby alarms, notifications, and changes are initiated relative to detected network congestion. The various states tracked by the active session state machine may be referred to as performance states or congestion states.

[0088] FIG. 10 also illustrates the storage of static attributes stored in memory and included within CPE **112**. The static attributes may include an access line or trunk throughput capacity, a determined unacceptable frame loss threshold, available CODEC capabilities, and an indication of a user end point type, such as wireless, mobile, fixed, or IP secure remote access.

[0089] FIG. 10 also illustrates that network performance information may be used to determine a long term state, a recent performance state, and a short term or immediate performance state. Measurements of CIR (real time) and EIR (non real-time) may be used to determine maximum sustainable and average throughputs over a long period of time (to determine a long term performance state), over a recent time period (to determine a recent performance state), and over each time interval immediately after it is completed (to determine an immediate performance state). The latter immediate performance state may be used, for example, to give a user or call controller information about an active call or other network communication.

[0090] Frame/packet counters, network probes, packet sniffers, rate limiters, and other performance measurement devices or processes may all be used to accomplish various aspects of embodiments of the current invention, whether implemented within a network device such as a switch, router, or CPE, collocated with such a network device, as a standalone network appliance, as an ASIC, or using any other suitable combination of hardware and software.

[0091] As discussed above, the combination of one or more sources of network performance information can be used to determine information about a network that is valuable in further improving network performance and identifying network failures. More particularly, network performance information regarding frame loss, CIRs between endpoints or on a

core network link, measured throughput of actual use, packet counters, determined frame sizes, and line rates may be used alone or in various combinations as valuable indicators of network performance and network failures. Such indicators may be used to determine compliance with service levels, identify causality for network failures, throttle permitted throughput rates, identify links having the lowest throughput capacity or effective throughput, or even alter how data is packetized, routed, and communicated over a network. In particular, as described in U.S. patent application Ser. No. 11/809,403, entitled System and Method for Establishing Calls Over a Call Path Having Best Path Metrics, which is hereby incorporated by reference, network performance information may be used to select CODECs for audio (including streaming audio and telephone calls) or video (including streaming video and video conferencing) data communications, to control CAC engines, for VoIP or other digital telephony call set-up, routing, throttling, provide indications of line state, or any other suitable selection, modification, or determination associated with network performance or the communication of data over one or more networks. As described in U.S. patent application Ser. No. 11/809,403, such network performance information can be collected or stored at individual network nodes or communicated over the network using packets such as OAM packets or other performance information packets. In one embodiment, as described, the network performance information can be communicated to a central resource used by one or more networks for different functions like call routing, determination of network egress points, least congested path determination, most available bandwidth determination, or any other suitable use. In various embodiments, the determinations described herein can be implemented by a session controller, session border controller, call controller, call control manager, network interface, CPE, or any other suitable device. For example, in one embodiment, session/call control logic illustrated as part of CPE 112 of FIG. 10 may accomplish the previously described functionality.

[0092] The previous detailed description is of a small number of embodiments for implementing the invention and is not intended to be limiting in scope. The following claims set forth a number of the embodiments of the invention disclosed with greater particularity.

What is claimed:

1. A method of determining the state of a path in a network, the method comprising:
 - measuring at a node of the network a rate of throughput over an interval for at least a portion of the path;
 - determining at the node of the network a derivation associated with a characteristic of the frames communicated over at least the portion of the path; and
 - determining a state of congestion of the path in response to the measured rate of throughput and the determined associated derivation.
2. The method of claim 1, wherein the associated derivation is an average packet size.
3. The method of claim 1, wherein the associated derivation is an amount of non-payload bits of one or more frames.
4. The method of claim 1, wherein the associated derivation is an average amount of non-payload bits of one or more frames.
5. The method of claim 1, wherein the associated derivation is an average amount of non-payload bits of frames received during a time interval.
6. The method of claim 1, wherein the associated derivation is a dynamic bandwidth profile.

7. The method of claim 1, wherein the associated derivation is a measured overall throughput less a determined amount of non-payload bits.

8. The method of claim 1, wherein the state of congestion is used to control traffic over the network.

9. The method of claim 1, wherein the state of congestion is used by a call control manager to route telephone calls.

10. The method of claim 1, wherein the state of congestion is used by a call control manager to perform call set-up for telephone calls.

11. A session controller comprising:

a memory operable to store the throughput capacity of a path in a network; and

a state machine in communication with the memory and operable to receive a measurement of a rate of throughput over an interval for the path, the state machine further operable to receive a derivation associated with a characteristic of frames communicated over the path;

wherein the state machine is further operable to determine a state of congestion of the path in response to the measurement, the throughput capacity, and the associated derivation.

12. The session controller of claim 11, and further comprising session control logic, the session control logic operable to modify at least one parameter of a communications session in response to the determined state of congestion.

13. The session controller of claim 11, and further comprising session control logic, the session control logic operable to modify the routing of communications in response to the determined state of congestion.

14. The session controller of claim 11, and further comprising session control logic, the session control logic operable to select a CODEC in response to the determined state of congestion.

15. The session controller of claim 11, and further comprising session control logic, the session control logic operable to select a network egress point in response to the determined state of congestion.

16. The session controller of claim 11, wherein the session controller is a call controller.

17. The session controller of claim 11, wherein the session controller is a call controller and the session control logic determines call set-up parameters.

18. The session controller of claim 11, wherein the session controller is a call controller and the session control logic determines a CODEC for a call.

19. The session controller of claim 11, wherein the session controller is a call controller and the session control logic determines call routing.

20. A computer-readable medium, the computer readable medium encoded with instructions operable when executed by a processor to:

measure at a node of the network a rate of throughput over an interval for at least a portion of the path;

determine at the node of the network a derivation associated with a characteristic of the frames communicated over at least the portion of the path; and

determine a state of congestion of the path in response to the measured rate of throughput and the determined associated derivation.

* * * * *